Customer Quality Newsletter

Contents

High Availability of Servers: Does number of nines mean anything?

Due to the explosive growth in the last decade of e-commerce and other computer-based services, a demand for systems to provide non-stop services has emerged. Customers rely so heavily on "mission" or "revenue" critical applications that any downtime, planned or unplanned, can be very expensive in terms of hard dollars, lost productivity or missed opportunities. To meet these requirements, systems must be designed such that they are highly available (HA).

What is Number of Nines?

A server's availability is expressed in terms of the *number of nines*. For example, 99.999 percent uptime or availability is referred to as five nines and 99.99 percent uptime is referred to as four nines. Five nines equates to 5.26 minutes downtime in a year assuming there are 365.25 days in a year. A common definition used for availability or uptime is the percentage of time in a year that a system is available to perform a specific function. What is not clear is how availability is measured or estimated for a system. If the system is down for maintenance, is it considered part of the downtime? If the system's performance is half its peak performance, is the system still considered as up or available? Suppose there are two systems with annual downtime of 5 minutes for each system. One system goes down for 5 minutes once in a year and the second system goes down five times in a year with each downtime lasting 1 minute. Are these systems equivalent in terms of availability? Are there ways of differentiating such systems?

As there is no industry-wide standard for measuring or estimating availability, it becomes very difficult to rely on the availability numbers unless the methodology for estimating availability is clearly specified. To ensure an accurate, consistent, and meaningful way of specifying availability, we hope the industry will adopt a methodology for specifying the *number of nines* to make it more meaningful and consistent.

How to Improve Availability

There are two main parameters that are important to system's availability – MTBF (mean time between failure) and MTTR (mean time to repair) as shown below.

$$Availability_{System} = \frac{MTBF_{System}}{MTBF_{System} + MTTR_{System}}$$

The MTBF is the expected time between component failures in a system. MTBF determines how often the system will experience a component failure though the component failure may not result in a system failure depending on component redundancy. As the system's MTBF is dependent mainly on redundancy and the quality of the components, increasing MTBF beyond a certain level can be a costly way of increasing availability as it may increase the cost of ownership drastically. In many cases, reducing MTTR may be a more cost effective way of increasing availability. This is why features such as fault management and prevention that reduce MTTR in a system are gaining more popularity.

A given HA objective for a system could be met with a relatively lower MTBF and very small MTTR as shown in Figure 1. Low MTBF implies frequent repairs, therefore more frequent visits to the site of the system, resulting in a higher maintenance cost. It may also require higher human intervention resulting in higher probability of human error. Though availability measures uptime of the system, it does not reflect on how many times the system will be required to be repaired and exposed to human intervention. This necessitates the right balance between MTBF and MTTR to meet a given HA objective and reduce maintenance cost.

How Intel Products Support HA

To ensure that the right balance between MTBF and MTTR is achieved for server systems to meet HA requirements, the main areas of focus are:

- Reduce component failures via maximizing

(continued on page 3)

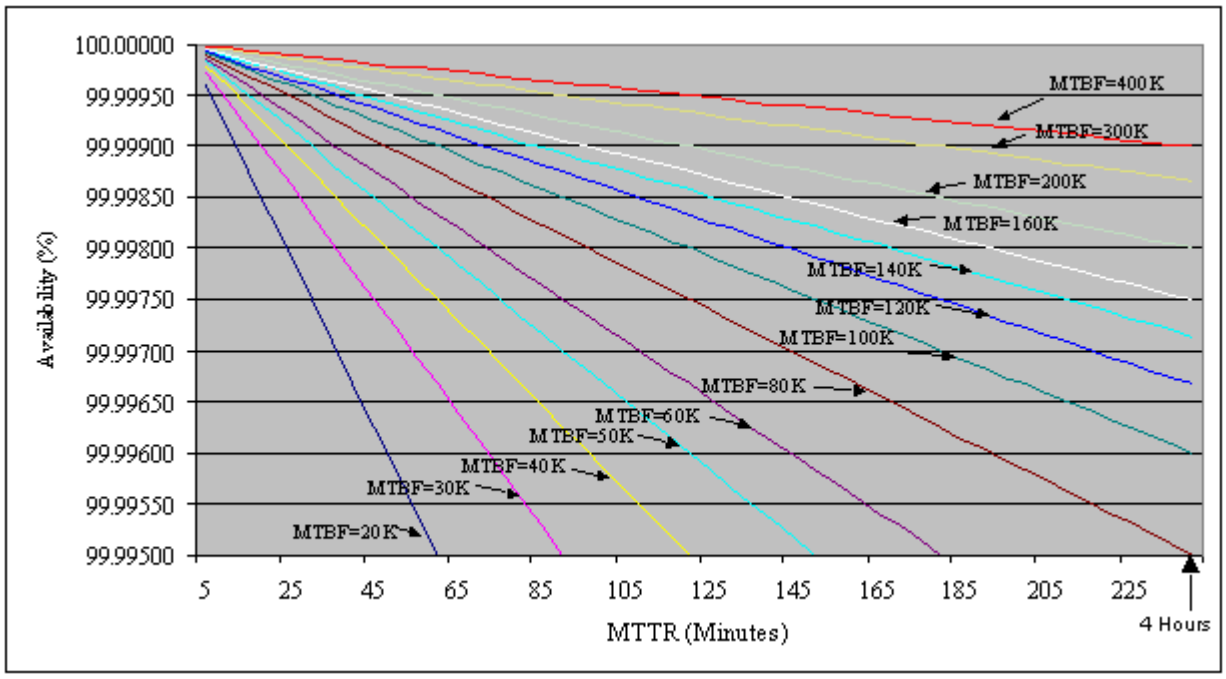

Figure 1. Availability versus MTTR

component's MTBF.

- Prevent system failures via maximizing system's MTBF.
- Minimize downtime by effectively managing system failures via minimizing system's MTTR.

Reducing Component Failures

Reducing component failures implies increasing component's MTBF. Components used are highly reliable and have high tolerance to environmental fluctuations such as power, temperature, vibration, etc. (NEBS Level 3 compliant). Components and systems are designed and built to ensure that they work well within their capacity. Factory testing strategy emphasizes on always meeting and using international standards and testing thresholds. System testing ensures that all components in a system work in harmony. Error checking and correction (ECC) memory and cache are supported on all Intel single board computers (SBCs).

Preventing System Failures

In a HA system, component failures are tolerable as long as a component failure does not result in a system failure. To minimize system failures, the system should support hardware redundancy such that the system has no single point of failures. To this goal, Intel's HA systems support redundant components such as power supplies, fans, data paths, switches, chassis management modules, connections to power outlet, etc. Thus, single points of failure are minimized if not eliminated. As backplane redundancy is costly, most of the current backplanes are passive and therefore have very high MTBF to reduce probability of failures.

Segmented backplanes are also currently being implemented whereby failure of one of the segments does not impact other segments. Thus, components connected to the functioning segments continue to provide service. This feature is very attractive in a telecommunication environment where complete system failure is considered to be catastrophic. Several of the products support RAID storage to minimize data loss. For master/slave architectures where a single master may control several peripherals, several of the Compact PCI products support redundant system slot (RSS) system masters (Active-Active and Active-Standby) to eliminate the single point of failure in the master.

One of the crucial ways of preventing system failure is by constantly monitoring and alerting for potential failures, overloads, component errors and components operating beyond their limits. These monitoring and alerting features

(continued on page 4)

that are common in Intel's highly available components and systems can reduce probability of system failure if corrective action is taken at the right time. These features include monitoring for temperature (e.g. CPU, boards, fan trays, power supplies, chassis, inlet and outlet air temperature), voltages (e.g. SBC, power supplies) and rpm (fans) and alerting with alarm levels – critical, major and minor.

Minimizing Downtime By Effectively Managing System Failures

One of the cost effective ways of reducing MTTR is by building management features into the system. Management features refer to features that help in fault detection and location; diagnosis; monitoring; reporting; and taking some corrective actions. Most of the management features contribute to availability in the following ways by:

- Reducing MTTR and human intervention
- Providing assistance when human intervention is required therefore reducing chances of human error.

The HA systems at Intel support several features that reduce MTTR and human intervention such as: (1) Capability to monitor system status, run diagnosis and take corrective actions such as soft and hard reset for components and the system, remotely; (2) Dynamic reconfiguration – Hot Swap Capabilities. When major components fail, it is important that the failed components are replaced and the server returned to full redundant status as soon as possible. The use of Hot Swap Capabilities provide for this and thereby maximizes service uptime; (3) Watchdog timers on components help in detecting failures earlier on; (4) BIOS recovery mode in SBCs help in reducing length of downtimes; (5) Auto restart where the system monitors the operating system whether it is alive and running. As soon as the operating system hangs for any reason, it automatically reboots the service and restarts the key applications; (6) Systems are also laid out to make it easy to remove and access components and simplify cabling.

To provide assistance for human intervention, LEDs indicate which components are active, failed, and in standby. Hot Swap Capabilities simplify maintenance and repair procedures.

Higher Number of Nines

To quench customers' thirst for higher *number of nines*, efforts continue in investigating several other intelligent features such as: (1) Thermal management of the system whereby the system has the capability to automatically increase or decrease fan speed, turn boards on or off, and modulate clock speed in response to system overheating while maintaining optimum system performance; (2) Power management whereby the chassis's overall power consumption is dynamically controlled by modulating clock speed or by powering down low priority components in the system while maintaining optimum system performance; (3) Failure prediction at component and system levels; (4) Support evaluation of errors using expert system techniques to carry out root cause analysis and provide effective repair strategy.

Several subset of these features are currently being investigated such as multi-speed fans for cooling, failure prediction using thresholds, capability to power down/up components in the system, etc. Though the wish list of HA features seems to be endless, we continue to explore and assess the feasibility of the features.

Breaking News

First Intel Quality-Focused Road Show in Brazil Touts Success

Intel hosted the first Brazil quality-focused road show in Sao Paulo, with more than 100 attendees from over 40 companies in attendance. During the December 2001 event, Intel managers from various internal business groups delivered presentations providing their expertise on each of the scheduled topics:

- Quality Operation Systems
- Virtual Factory and Copy Exactly
- Product Qualification Methodology
- The Intel® Pentium® 4 processor Manufacturing Quality
- Intel's Lead Free Program

To facilitate presentations, all materials were translated to Portuguese with simultaneous translation available to non-English speaking attendees.

This event was an excellent venue for sharing Intel's manufacturing quality leadership with local customers, thus increasing customer awareness of Intel's customer quality

(continued on page 5)

Breaking News (continued)

support initiatives. The activities focused on:

- Delivering a better understanding of Intel's quality strength through sharing Intel's quality operating system & methodology with valued customers
- Enhancing the relationship of Intel and our customers through the sharing of best manufacturing quality practices
- Demonstrating Intel's efforts by delivering quality support and technology leadership to the industry
- Making Intel's service quality a key and positive value in doing business with Intel

Brazil sets the stage for information sharing. Intel managers and local customers met in Sao Paulo to exchange ideas and focus on quality support initiatives.

The audience returned positive feedback and also included topics they would like to see in future events of this kind. Another road show is in the plans for the first half of 2002.

Do You Know?

Reinforcing Manufacturing Leadership Through Intel's Quality Management System

Intel has established robust quality systems that enable our high volume manufacturing sites to achieve and maintain world-class quality levels. Intel's quality management

Do You Know? (continued)

system encompasses the processes defined from receiving the silicon from our suppliers, through manufacturing, to shipping the right products on time to our end customers. Currently, all of Intel's manufacturing sites (fabrication and assembly test) are registered to the ISO 9001/9002 International Standard as the baseline for our quality management system.

As a requirement of both the ISO9001/9002 standard and Intel, a great deal of effort is placed into constantly improving our processes and systems. There are several monitors in place that ensure all quality and reliability requirements have been met, and that opportunities for continual improvement have been examined. Some examples include

- an internal quality audit system that looks at the effectiveness of our processes and systems,
- reliability monitors to measure reliability performance on semi-finished or finished products
- and customer monitors to ensure that our products perform according to our customer's environment.

Through these various monitors, we can identify nonconformance issues, and track the appropriate corrective and preventive actions to a successful closure. Nonconformance issues are driven to closure by

- containing all affected material,
- conducting a risk assessment to develop an evaluation plan,
- identifying and implementing a corrective and preventive action,
- and the proper disposition of affected material.

For continual improvement and to prevent repeat issues, data is gathered, reviewed, and trended to make fundamental improvements.

Intel's quality management system consists of several elements resulting in a closed loop and systematic process. The methodology reinforces Intel's leadership position by providing high-level requirements to achieve synergy, as well as operational procedures to drive consistent execution. Training on the methodology and supporting

(continued on page 6)

Do You Know? (continued)

tools is maintained and delivered to all personnel. Indicators are gathered, trended, and analyzed on a periodic basis. Feedback mechanisms exist to gather input from both internal and external customers. The data and feedback drive continual improvement of our quality management system.

Intel's Quality Principle

World-class quality is an essential ingredient for Intel's success. In order to maintain our industry leadership, we are committed to making quality an integral part of every aspect of our business. As a diverse company serving numerous markets and customers, we will also continuously strive to deliver an outstanding quality image for all Intel and affiliated brands.

Policy Statement

All Intel businesses and subsidiaries will be recognized as #1, #2 or excellent for product and service quality in the markets they serve.

Customer Feedback

Opportunity to Share Your Thoughts

We value your readership and would like to hear from you. Please send your feedback to the Intel Customer Quality Newsletter e-mail address: customer.quality.newsletter@intel.com

Quality Information Center

FACR:	1-800-628-8686 (in North America) or contact your local Customer Quality Engineer
Intel Developer:	developer.intel.com
Intel's OEM Customer:	oem.intel.com
Intel's Channel Support:	channel.intel.com
Contact your local sales representative for more information or visit our web site at:	support.intel.com
Feedback and Comments:	
Contact your local Field Sales Engineer or Customer Quality Engineer for more details.	
Send us an e-mail:	customer.quality.newsletter@intel.com